arXiv:1512.08168v5 [cs.FL] 11 Feb 2016

Quick Brown Fox in Formal Languages

Kazuhiro Inaba (kiki@kmonos.net)

February 12, 2016

Abstract

Given a finite alphabet Σ and a deterministic finite automaton on Σ, the problem of determining whether the language recognized by the automaton contains any *pangram* is NP-complete. Various other language classes and problems around pangrams are analyzed.

1 Introduction

Definition 1.1. A string $w \in \Sigma^*$ is a *pangram* (*perfect pangram*, respectively) on finite alphabet Σ if all $\sigma \in \Sigma$ appear at least once (exactly once) in w. We denote the set of all pangrams on Σ by P_Σ and the set of all perfect pangrams by E_Σ.

The famous English pangram "The quick brown fox jumps over the lazy dog" is in P_Σ for $\Sigma = \{\mathtt{a}, \ldots, \mathtt{z}\}$, while it is not in E_Σ because some of the letters appear more than once. In Japanese, there is an old poem called Iroha-Uta[1] being a perfect pangram over 47 letters of Japanese alphabet. In the programming language Ruby, a perfect pangram on ASCII printable characters can be a valid program that outputs each ASCII printable character exactly once.[2]

No one will doubt that coming up with the artistic sentence of the "quick brown fox" or Iroha-Uta is a tough brain teaser. The purpose of this paper is to confirm the natural intuition; creating a pangram is indeed hard. We discuss on the computational complexity of the problem deciding whether there is a (perfect) pangram in a given language.

2 Problem Definitions

We assume the reader to be familiar with basic concepts of formal languages and computational complexity, e.g., by [RS97, GJ79].

[1] https://en.wikipedia.org/wiki/Iroha
[2] https://github.com/tric/trick2013/tree/master/kinaba

Definition 2.1. A *language acceptor class* X consists of a sets X_Σ of *language acceptors* for each finite alphabet Σ and a membership relation $mem_X \in \bigcup_\Sigma X_\Sigma \times \Sigma^*$.

Examples of language acceptor class are DFA (deterministic finite automata), NFA (deterministic finite automata), CFG (context free grammars), etc. Each member of the set $\mathsf{DFA}_{\{0,1\}}$ is a deterministic finite automaton over the binary alphabet. We assume that each member $x \in X_\Sigma$ has a naturally associated size $|x|$, such as the number of transitions or grammar rules. For $x \in X_\Sigma$, we write $\mathcal{L}(x)$ for $\{w \in \Sigma^* \mid mem_X(x, w)\}$.

Definition 2.2 (Pangram Problem)**.** *Pangram problem* (respectively, *perfect pangram problem*) for a language acceptor class X is, given a finite alphabet Σ and a language acceptor $x \in X_\Sigma$ as an input, to ask whether $\mathcal{L}(x) \cap P_\Sigma$ (resp. $\mathcal{L}(x) \cap E_\Sigma$) contains some member or not.

Informally speaking, the pangram problem gives us a grammar of a language and asks if there can be a pangram in the language. In the rest of the paper, we measure the computationally complexity of (perfect) pangram problems in terms of $|x| + |\Sigma|$.

Here are some easy facts on pangrams, from the state complexity perspective:

Proposition 2.3. P_Σ is a regular language. The minimum deterministic finite automata recognizing P_Σ has $2^{|\Sigma|}$ states.

Proposition 2.4. E_Σ is a regular language (actually a finite language.) The minimum deterministic finite automata recognizing E_Σ has $2^{|\Sigma|} + 1$ states.

Proof. Count the number of Myhill-Nerode equivalence classes. □

These exponential natures block us from giving efficient algorithms by a naive construction for many decision problems on pangrams, as we see in the following sections.

3 Hardness of Pangram Problems

Firstly, the problem is NP-hard already for DFA.

Theorem 3.1. Perfect pangram problem for DFA is NP-hard.

Proof. The proof is by reduction from Hamiltonian Path Problem, which asks, given a directed graph (V, E), if there is a simple path visiting all the nodes exactly once.

The instance of Hamiltonian Path Problem can be converted to a perfect pangram problem for DFA as follows. We let the alphabet $\Sigma = V$ and the set of states of the automaton to be $Q = V \cup \{q_{\mathrm{src}}, q_{\mathrm{fail}}\}$. The initial and the final

states of the automaton is q_{src} and $Q \setminus \{q_{\mathrm{fail}}\}$, respectively. We let the transition function

$$\begin{aligned}
\delta(q_{\mathrm{src}}, u) &= u \text{ for all } u \in V \\
\delta(v, u) &= u \text{ for all } v, u \in V \text{ where } (v, u) \in E \\
\delta(v, u) &= q_{\mathrm{fail}} \text{ for all } v, u \in V \text{ where } (v, u) \notin E \\
\delta(q_{\mathrm{fail}}, u) &= q_{\mathrm{fail}} \text{ for all } u \in V
\end{aligned}$$

Then, if there is a Hamiltonian path $v_1 \to \cdots \to v_n$, there exists a pangram $v_1 \cdots v_n$ in the language accepted by the automaton, and vice versa. □

Theorem 3.2. Pangram problem for DFA is NP-hard.

Proof. Reduction from the perfect pangram problem for DFA. Note that for each Σ, the set $S_\Sigma = \{w \mid |w| = |\Sigma|\}$ is a regular language and there exists a deterministic finite automaton s_Σ for the language, with size polynomial in Σ.

Then, given a problem instance (Σ, x) of the perfect pangram problem, $\mathcal{L}(x)$ contains a perfect pangram if and only if $\mathcal{L}(x) \cap \mathcal{L}(s_\Sigma)$ contains a pangram, since $\mathcal{L}(s_\Sigma)$ fixes the string length and makes all pangrams also perfect pangrams. Computing the deterministic finite automaton representing the the intersection is well-known to be done in polynomial time by the product construction. □

The automaton constructed in the proof of Theorem 3.2 represents a *finite* language. In other words, the NP-hardness arises even from such a restricted class of languages.

Proposition 3.3. Perfect pangram problem and pangram problem are NP-hard for FinDFA, where FinDFA is the subclass of DFA representing only finite languages.

Interestingly, a language whose complement is finite (called *cofinite* languages) also exhibits the same hardness on the perfect pangram problem. Note that the normal pangram problem for CofinDFA is trivial; a cofinite language always contains a pangram because there are infinitely many pangrams.

Proposition 3.4. Perfect pangram problem is NP-hard for CofinDFA, where CofinDFA is the subclass of DFA representing only cofinite languages.

Proof. Reduction from the perfect pangram problem of general DFA. Assume a DFA g is given, and let e_Σ be a DFA representing the set $\{w \mid |w| \neq |\Sigma|\}$, which can be constructed in $|\Sigma| + 1$ states. Then, $\mathcal{L}(g)$ contains a perfect pangram if and only if $\mathcal{L}(g) \cup \mathcal{L}(e_\Sigma)$ does. The latter is a cofinite language. □

Now, let us turn our eyes to the containment in NP, to complete the proof of the NP-completeness.

Theorem 3.5. Perfect pangram problem is in NP for CFG.

Proof. The witness, i.e. a perfect pangram string, has length Σ and can be guessed by a nondeterministic Turing machine in linear time. Then, since the combined complexity of checking the membership of the string against the acceptor is in polynomial time, the whole step of checking perfect pangram problem is in NP. □

The result extends to classes with higher expressiveness as long as they permit NP-time membership judgment, but we need some care. The membership check needs to be in NP in combined complexity, meaning that both the string and the language acceptor are counted as inputs.

For instance, although the membership problem for the higher order version of CFG known as IO/OI-hierarchy [Dam82] is in NP [IM08], the complexity is only with respect to the string size and a grammar is considered to be fixed. If we take the grammar as a part of input, the membership problem is EXP-complete [TK86] already for indexed languages [Aho68, Fis68], which is only one level above CFG. It remains PSPACE-complete even after limiting to the subclass of grammars without ϵ-rules [Ost15].

Compared to perfect pangrams, a pangram may have arbitrary length. It may not be easy to find a short witness directly. For instance, there exists a context-free grammar whose shortest member string is exponentially larger than the grammar. Nevertheless, the problem is proved to stay in NP by applying a simple grammar transformation.

For preparation, let's call a string u is a *subsequence* of a string w and write $u \sqsubseteq w$ if u is obtained by deleting several letters from w without changing the order of the remaining letters. The *downward closure* of a language L is the set of strings $L_{\downarrow} = \{u \mid u \sqsubseteq w \text{ for some } w \in L\}$. The important fact is that L contains a pangram if and only if $L_{\downarrow}$ contains a perfect pangram.

Theorem 3.6. Pangram problem is in NP for CFG.

Proof. Pangram problem reduces to perfect pangram problem as long as the downward closure for CFG can be constructed in polynomial time–and indeed it is. Downward closure of CFG can be computed in polynomial time, for instance by constructing Chomsky normal form and then adding a rule $N \to \epsilon$ for all nonterminals N. □

Putting altogether and considering that conversion from DFA to CFG is trivially done in polynomial time, we get the following summary of this section.

Corollary 3.7. Perfect pangram problem is NP-complete for FinDFA, CofinDFA, DFA, NFA, and CFG. Pangram problem is NP-complete for FinDFA, DFA, NFA, and CFG.

4 Pangram-Cover Problems

Viewing from the other side, asking if a language contains some pangram is equivalent to ask if the complement of the language contains all pangrams,

and to negate the answer. Thus we obtain the following yet another hardness statement.

Corollary 4.1. Given a language acceptor g, the problem asking whether $\mathcal{L}(g) \supseteq E_\Sigma$ is coNP-complete for FinDFA, CofinDFA, DFA, NFA, and CFG.

Proof. It is in coNP because the negative witness can be given by a non-member perfect pangram, which is in linear size. It is coNP-hard because the negation of this problem is equivalent to the perfect pangram problem for the complement of g. Since the perfect pangram problem is NP-hard for FinDFA, CofinDFA, and DFA, for the language acceptors that can subsume the polynomial-time constructible complement of either of the three classes, the result follows. □

The pangram counterpart of this "contains all" question becomes truly harder for NFA and CFG, which are not polynomial-time closed under complement.

Theorem 4.2. Given a language acceptor g, the problem asking whether $\mathcal{L}(g) \supseteq P_\Sigma$ is coNP-complete for CofinDFA and DFA, *undecidable* for CFG, and PSPACE-complete for NFA.

Proof. The coNP-hardness follows by the same argument as Corollary 4.1. The coNP containment for the CofinDFA and the DFA cases holds because for DFA we can still always find a short negative witness by a slight modification of the well-known *pumping lemma*.

The undecidability for CFG is proved by the reduction from the universality problem, which asks whether $\mathcal{L}(g) = \Sigma^*$ or not. This problem is known to be undecidable [GR63]. Now, assume a CFG g over alphabet $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ is given. Let $w = \sigma_1 \cdots \sigma_n$, $w\mathcal{L}(g)$ the language obtained by concatenating w in front of each string in $\mathcal{L}(g)$, and NS_w the set of strings that do *not* start with w. Then it is easy to verify:

$$\begin{aligned} & \mathcal{L}(g) = \Sigma^* \\ \iff & w\mathcal{L}(g) \cup NS_w = \Sigma^* \\ \iff & w\mathcal{L}(g) \cup NS_w \supseteq P_\Sigma. \end{aligned}$$

Since, NS_w is a context free language and CFG is closed under concatenation and union, the decision procedure of the last inequality will decide the universality problem as well, which is impossible.

The PSPACE-hardness for NFA is similarly derived from the universality problem (known to be PSPACE-complete for NFA [MS72, III73]) by using the fact that NS_w is representable in polynomial-size NFA and so for the union.

The containment in PSPACE can be proved by the backward reduction. Notice that the set $\overline{P_\Sigma}$ of non-pangrams can be represented by a polynomial-size NFA by seeing the language as $\bigcup_{a\in\Sigma}(\Sigma\backslash\{a\})*$. Thus, by testing the universality of $\mathcal{L}g \cup \overline{P_\Sigma}$ we can decide whether $\mathcal{L}g \supseteq P_\Sigma$ or not in PSPACE. □

5 On Strictly Locally Testable Languages

It should be natural to ask if we can find a weaker class of language acceptors that admits tractable (perfect) pangram problems. The problems were hard even for finite languages, but they are not the only sub-regular class.

Our first choice of the trial is *local languages* [Med64], or equivalently *2-strictly locally testable languages* [MP71, BS73]. For an integer $k \geq 2$, k-strictly locally testable languages (k-slt for short) are represented by three sets $S \subseteq \Sigma^{k-1}$, $I \subseteq \Sigma^k$, and $E \subseteq \Sigma^{k-1}$. A string w belongs to the k-slt language denoted by (S, I, E) if and only if w begins with a prefix in S, ends with a suffix in E, and all k consecutive substrings in w are in I. We call k-SLT the class of k-slt language acceptors whose size is measured by the sum of representing set sizes $|S| + |I| + |E|$.[3]

Unfortunately, the perfecet pangram problem still remains to be NP-hard.

Proposition 5.1. Perfect pangram problem is NP-hard for 2-SLT.

Proof. The same proof as Theorem 3.1 applies, because the represented language is actually 2-slt, with $S = E = \Sigma$ and edges denote the permitted neighbors I. □

Contrary to the perfect pangram problem, the normal pangram problem becomes tractable when it is limited to 2-SLT. The product construction proof of Theorem 3.2 does not apply to 2-SLT because the resulting new language is not in 2-slt anymore.

Theorem 5.2. Pangram problem for 2-SLT can be solved in linear time.

Proof. Let (S, I, E) be a tuple representing a local language, and consider a graph with Σ is the set of nodes and I the set of edges. Whether or not a pangram is contained in the language is equivalent to ask whether there is a (not necessarily simple) path beginning from S and ending in E that visits all nodes Σ.

This graph-theoretic problem is solvable through the decomposition to strongly-connected components. A path visiting all nodes in the original graph exists if and only if the acyclic graph obtained by contracting strongly-connected components has a Hamiltonian path from a component containing a S node to a component with E. For an acyclic graph, the Hamiltonian path (if any) is uniquely obtained by topologically ordering the nodes. All these condition and checkable in linear time. □

If we go one step higher to the class of 3-SLT, the pangram problem becomes hard again.

Theorem 5.3. Pangram problem is NP-hard for 3-SLT.

[3] Since we can convert this representation to a minimal DFA in polynomial time, the NP-hardness results in this section still hold even if we assume the k-SLT is given by a DFA.

Proof. Again the proof is by reduction from Hamiltonian Path Problem. Let us assume the given instance of Hamiltonian Path Problem is (V, E). We let the alphabet $\Sigma = V \cup \{1, \ldots, |V|\}$, and the corresponding instance (S_l, I_l, E_l) of 3-SLT is constructed as follows:

$$\begin{aligned} S_l &= \{(v, 1) \mid v \in V\} \\ I_l &= \{(v, k, u) \mid (v, u) \in E, k \in \{1, \ldots, |V|\}\} \\ &\quad \cup \{(k, v, k+1) \mid v \in V, k \in \{1, \ldots, |V| - 1\}\} \\ E_l &= \{(v, |V|) \mid v \in V\} \end{aligned}$$

Then, if there is a Hamiltonian path $v_1 \to \cdots \to v_n$ there exists a pangram $v_1 1 v_2 2 \cdots (n-1) v_n n$ in the language accepted by the automaton, and vice versa. □

6 On Strictly Piecewise Testable Languages

Beside the hierarchy of locally testable languages, yet another well-studied classes of sub-regular languages is those of *piecewise testable languages* [Sim75]. Recall that by $u \sqsubseteq w$ we mean u to be a (not necessary contiguous) subsequence of w. A language L is *k-strictly piecewise* [Hei07, RHB+10] if L can be defined in terms of forbidden subsequences, i.e., if L can be written as $L = \{w \mid \forall u \in F. u \not\sqsubseteq w\}$ for some $F \subseteq \Sigma^{\leq k}$. Let k-SPT be the class of k-strictly piecewise language acceptors whose size is measured by the size $|F|$ of the forbidden subsequence set.

First of all, since all strictly piecewise languages are downward closed, it contains a pangram if and only if it contains a perfect pangram. Hence the two problems become the same. For a lower class of this hierarchy, the problems become tractable.

Proposition 6.1. Perfect pangram problem and pangram problem can be solved in linear time for 2-SPT.

Proof. If F contains an empty string or a string of length-1, the language trivially excludes pangrams. Hence we assume $F \subseteq \Sigma^2$ below.

Note that each member (x, y) in F represents the constraint that the letter x must come *after* y in a pangram (otherwise x comes before y in the pangram, which contradicts F.)

Now, consider a graph with nodes Σ and edges F. If this graph has a cycle, no pangram is contained in the language because all pangram must violate the "must come after" constraint for some edge in the cycle. If this graph has no cycle, then the reverse topological sort will give a pangram, in which all x comes after y for all $(x, y) \in F$ by the definition of topological sorting. Cycle detection of a directed graph can be done in linear time. □

Not surprisingly, once we change the parameter from 2 to 3, the NP-hardness creeps in.

Theorem 6.2. Perfect pangram problem and pangram problem are NP-complete for 3-SPT.

Proof. It is in NP because the witness of perfect pangram is in linear size.

The reduction for NP-hardness is from the Betweenness Problem [Opa79], that asks there is a total ordering of elements of a finite set Σ, under the set of given constraints $(a, b, c) \in \Sigma^3$ that imposes either $a < b < c$ or $c < b < a$ (i.e., b must be in between the other two.)

We can construct the forbidden set of 3-SPT over the same Σ. For each betweenness constraint (a, b, c), the forbidden subsequences "acb", "cab", "bac", and "bca" are added. Then, if a total ordering exists, the element of Σ listed along the order gives a perfect pangram over Σ avoiding all forbidden subsequences, and vice versa. □

Interesting open problem is the complexity when the 3-SPT language is specified in the form of a DFA. The minimum DFA representation can be exponentially larger than the forbidden set representation, hence our hardness result does not directly apply.

7 Tractability of Always-Pangram Problems

Another interesting question on a language acceptor is to decide whether all the strings it accepts are (perfect) pangrams. Despite the NP-hardness of showing existence of one pangram, checking if *all* of them are pangram is efficiently decidable.

Theorem 7.1. For a CFG g over alphabet Σ, whether or not $\mathcal{L}(g) \subseteq P_\Sigma$ is decidable in $O(|\Sigma| \cdot |g|)$ time.

Proof. For each $\sigma \in \Sigma$, compute the intersection of g and $(\Sigma \setminus \{\sigma\})^*$. This is obtained by just dropping all the rules containing σ in right-hand side. If the obtained grammar recognizes a non-empty language (which is linear time decidable), there is a non-pangram member in $\mathcal{L}(g)$. If all of the intersections are empty, $\mathcal{L}(g) \subseteq P_\Sigma$ holds. □

Corollary 7.2. For a CFG g over alphabet Σ, whether or not $\mathcal{L}(g) \subseteq E_\Sigma$ is decidable in $O(|\Sigma| \cdot |g|)$ time.

Proof. $\mathcal{L}(g) \subseteq E_\Sigma$ is equivalent to saying that $\mathcal{L}(g) \subseteq P_\Sigma$ and all members of $\mathcal{L}(g)$ is of length $|\Sigma|$. The latter condition is easily checkable in $O(|g|)$ time, by assigning the length for each nonterminal in a bottom-up manner (from the ones with right-hand side consisting of terminals) and checking no contradiction. □

The problems analyzed in this section have natural interpretation when the language is considered to be describing some sequence of events possibly happening. The condition $\mathcal{L}(g) \subseteq P_\Sigma$ corresponds to say that all events will eventually happen for all possible event sequences.

	$\mathcal{L} \cap P_\Sigma \neq \emptyset$	$\mathcal{L} \cap E_\Sigma \neq \emptyset$	$\mathcal{L} \supseteq P_\Sigma$	$\mathcal{L} \supseteq E_\Sigma$	$\mathcal{L} \subseteq P_\Sigma$	$\mathcal{L} \subseteq E_\Sigma$
2-SLT	P	NPc	-	-	P	P
3-SLT	NPc	NPc	-	-	P	P
2-SPT	P	P	-	-	-	-
3-SPT	NPc	NPc	-	-	-	-
CofinDFA	-	NPc	coNPc	coNPc	-	-
FinDFA	NPc	NPc	-	coNPc	P	P
DFA	NPc	NPc	coNPc	coNPc	P	P
NFA	NPc	NPc	PSPACEc	coNPc	P	P
CFG	NPc	NPc	undecidable	coNPc	P	P

Table 1: Summary of the results (- means the problem is almost trivial.)

8 Related Work

The All Colors Shortest Path problem (ACSP) [BÇG$^+$15] asks the shortest path in an undirected graph with nodes colored, under the constraint that the paths must visit all colors. The pangram problem for NFA can be regarded as an edge-colored and directed version of All Colors Path problem. Our result shows that the problem is NP-complete even if the shortestness condition is dropped and just the existence of such paths is asked.

9 Conclusion

We have defined the notion of pangram and perfect pangram in terms of formal language theory. Computational complexity of several problems around pangrams are investigated. The result presented in this paper is summarized in Table 1.

References

[Aho68] Alfred V. Aho. Indexed grammars–an extension of context-free grammars. *Journal of the ACM*, 15:647–671, 1968.

[BÇG$^+$15] Yunus Can Bilge, Dogukan Çagatay, Begüm Genç, Mecit Sari, Hüseyin Akcan, and Cem Evrendilek. All colors shortest path problem. `http://arxiv.org/abs/1507.06865`, 2015.

[BS73] J.A. Brzozowski and Imre Simon. Characterizations of locally testable events. *Discrete Mathematics*, 4:243–271, 1973.

[Dam82] Werner Damm. The IO- and OI-hierarchies. *Theoretical Computer Science*, 20:95–207, 1982.

[Fis68] Michael J. Fischer. Grammars with macro-like productions. In *Switching and Automata Theory*, pages 131–142, 1968.

[GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.

[GR63] Seymour Ginsburg and Gene F. Rose. Some recursively unsolvable problems in ALGOL-like languages. *Journal of the ACM*, 10:29–47, 1963.

[Hei07] Jeffrey Heinz. *The Inductive Learning of Phonotactic Patterns*. PhD thesis, University of California, Los Angeles, 2007.

[III73] H. B. Hunt III. On the time and tape complexity of languages I. In *ACM Symposium on Theory of Computing (STOC)*, 1973.

[IM08] Kazuhiro Inaba and Sebastian Maneth. The complexity of tree transducer output languages. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 244–255, 2008.

[Med64] Yu T. Medvedev. On the class of events representable in a finite automaton (English translation). In *Sequential Machines: Selected Papers*. Addison-Wesley, 1964.

[MP71] Robert McNaughton and Seymour A. Papert. *Counter-Free Automata (M.I.T. Research Monograph No. 65)*. The MIT Press, 1971.

[MS72] A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Switching and Automata Theory*, pages 125–129, 1972.

[Opa79] J. Opatrny. Total ordering problem. *SIAM Journal on Computing*, 8:111–114, 1979.

[Ost15] Johannes Osterholzer. Complexity of uniform membership of context-free tree grammars. In *International Conference on Algebraic Informatics (CAI)*, pages 176–188, 2015.

[RHB+10] James Rogers, Jeffrey Heinz, Gil Bailey, Matt Edlefsen, Molly Visscher, David Wellcome, and Sean Wibel. On languages piecewise testable in the strict sense. In *The Mathematics of Language (MOL)*, pages 255–265, 2010.

[RS97] Grzegorz Rozenberg and Arto Salomaa, editors. *Handbook of Formal Languages, Vol. 1: Word, Language, Grammar*. Springer-Verlag New York, Inc., New York, NY, USA, 1997.

[Sim75] Imre Simon. Piecewise testable events. In *Automata Theory and Formal Languages*, pages 214–222, 1975.

[TK86] Shinichi Tanaka and Takumi Kasai. The emptiness problem for indexed language is exponential-time complete. *Systems and Computers in Japan*, 17:29–37, 1986.